

Non-parametric Error Modeling for Ultra-wideband Localization Networks

Acshi Haggemiller

Maximilian Krogius

Edwin Olson

Abstract—We propose an ultra-wideband-based (UWB) localization system that achieves high accuracy through non-parametric estimation of measurement probability densities and explicit modeling of antenna delays. This problem is difficult because non-line-of-sight conditions give rise to multimodal errors, which make linear estimation methods ineffective. The primary contribution in this paper is an approach for both characterizing these errors in situ and an optimization framework that recovers both positions and antenna delays. We evaluate our system with a network of 8 nodes based on the DecaWave DWM1000 and achieve accuracies from 3 cm RMSE in line-of-sight conditions to 30 cm RMSE in non-line-of-sight conditions. Collecting measurements and localizing the network in this manner requires less than a minute, after which the realized network may be used for dynamic real-time tracking.

I. INTRODUCTION

Wireless sensor networks show promise for rapid, accurate, and inexpensive deployment of localization in new environments, especially compared to sensors such as LIDAR or motion capture. We use ultra-wideband (UWB) radios because of their low cost and potential for high accuracy measurements; specifically, we use the DecaWave DWM1000. A key problem getting high accuracy with this device is the characterization of an internal UWB parameter called *antenna delay*, which is the combined time it takes a signal to leave the generating hardware and be radiated by the antenna and the similar delays on the receiver. For the DWM1000, single-measurement error improves from roughly 30 cm to 4 cm after the process of manual antenna delay calibration [1]. Another problem is that the localization accuracy of a robot or other moving UWB node depends on knowing the positions of the fixed nodes, but manually measuring these for a large network can be time consuming.

In order to achieve high accuracy and support rapid deployment, we consider the joint problem of solving for both antenna delays and node positions, using only the direct *time of propagation* measurements (sum of the respective antenna delays and time-of-flight) between nodes in the network.

This approach is difficult for several reasons. First, the addition of antenna delays and lack of a priori known *anchors* makes the problem less constrained. Second, in non-line-of-sight (NLOS) conditions the signal strength for the direct path is weakened and the timed signal will more likely correspond to a non-direct reflected route, or *multipath*, that will overestimate the true distance. Third, to compensate for the weaker signal strength of the true path signal, the

The authors are with the Computer Science and Engineering Department, University of Michigan, Ann Arbor, MI 48104, USA. {acshikh,mkrogius,ebolson}@umich.edu

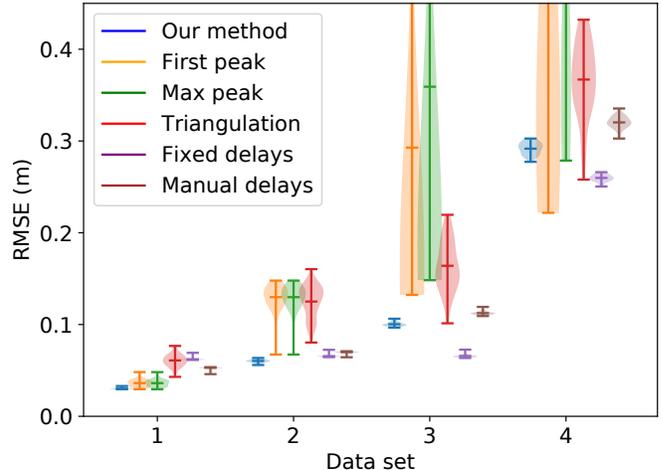


Fig. 1. Robustness of our method to increasingly non-line-of-sight (NLOS) datasets. We show the distribution of RMSE of 7 UWB node locations evaluated on 50 bootstrap samples from each dataset. First peak and max peak methods, as predicted, are not robust to NLOS conditions (datasets 2, 3, 4). Triangulation, which does not perform non-linear optimization, has higher RMSE. Fixed nominal and manual delay methods cannot achieve as high accuracies on the low-NLOS datasets 1 and 2 as the full method, and the manual values might be over-fit to calibration conditions. High uncertainty in datasets 3 and 4 cause the full method to perform worse than the constant delay methods because the system is less constrained.

UWB radio can be made more sensitive, but this increases the likelihood of mistaking noise for the path signal and underestimating the true distance [2]. The combination of the previous effects leads to a measurement error model dependent on unknown environment conditions with an unknown number of modes and possible non-Gaussian noise. Finally, the optimization problem is non-convex, has many local minima, and can suffer from geometric flip ambiguities [3].

To solve for unknown antenna delays, we explicitly model these delays and their prior probabilities. To account for unknown error conditions, we directly approximate the measurement probability density function (PDF) in a non-parametric fashion. As the problem is non-convex, we use gradientless non-linear optimization initialized with trilateration through robust quadrilaterals [4].

The main contributions of this paper include:

- Explicit coupled modeling of node antenna delays and locations in a maximum a posteriori formulation.
- A framework for incorporating non-parametric error models, where the error models are built dynamically based on empirical data for each pairwise radio link.
- Four open-access datasets which correspond to varying levels of LOS and NLOS links, for comparisons with our method.

II. RELATED WORK

The use of sensor networks for localization is well studied in literature. A variety of sensor types may be used, including those based on WiFi access points [5], Bluetooth Low Energy (BLE) devices [6], chirp spread spectrum (CSS) radios [7], and ultra-wideband (UWB) radios [8]. While WiFi and BLE-based methods often work opportunistically with existing hardware, their dependence on received signal strength (RSS) and fingerprinting metrics leads to relatively coarse location estimates (on the order of meters). Furthermore, because these techniques do not directly use range estimates, they require a prior map of the localization environment.

When the node sensors are capable of taking range measurements, the main problem is *network localization*, or the localization of unknown node positions given the positions of a priori known anchors.

Three existing methods for network localization are: simulated annealing, semidefinite programming, and multidimensional scaling. When using simulated annealing localization [9], the node position estimates are randomly perturbed by progressively decreasing amounts which correspond to a temperature parameter. Changes that lower the error function are more likely to remain. At each temperature, all the nodes are perturbed sufficiently for the system to enter a thermal equilibrium before the temperature is lowered. When using semidefinite programming [10], a suboptimal relaxation of the weighted maximum likelihood estimate is used. This SDP formulation is convex and provides a good initial estimate for further refinement through a gradient search method. When using multidimensional scaling [11], the objective function is the least squares error between measured and expected ranges. A gradient estimate is found that is exact when the error is zero, which is then used with gradient descent. Unlike our method, each of these uses known anchors and was evaluated only in simulation, not taking into account hardware complexities like antenna delay.

Graph realization is a subproblem of network localization where all node positions are solved for with only range measurements. Known anchors are not used. This subproblem is less common in the literature. Di Franco et al. [12] propose an MDS-based method with Gaussian mixtures and expectation-maximization to solve the graph realization. They evaluate their method with the SurePoint [13] UWB hardware, which is also based on the DecaWave DW1000 UWB chip. They do not consider antenna delay characterization of their hardware.

Even with the use of UWB radios, there are multiple ways to estimate range between devices. In one case, by using time difference of arrival and wired synchronized clocks, a UWB system achieved sub-millimeter accuracy [8]; however, this need for wired time synchronization could make the system infeasible in certain scenarios. In a more typical example of synchronization-free time-of-flight (TOF) ranging, a UWB system also based on the DWM1000 module achieved single measurement error in the range of about 12 cm to 36 cm [14].

Once the graph realization is complete, the fixed UWB

nodes can be used for real-time tracking of dynamic objects. Prior works in this area generally perform sensor fusion to mitigate multipath effects and achieve higher update rates. Hol et al. fuse inertial measurement unit (IMU) readings with UWB time-of-arrival measurements from six synchronized nodes [15]. Zwirello et al. assume a pedestrian model and also fuse step length estimates in addition to IMU [16].

The area of graph realization for UWB networks has not yet been extensively explored. Our key differences compared to prior work include modeling of antenna delays and optimization on an empirical error model.

III. APPROACH

This paper discusses use of UWB-based nodes using standard double-sided two-way ranging (DS-TWR) time-of-flight (TOF) measurements with the inexpensive DecaWave DWM1000 UWB module. We also constrain the problem to two dimensions because of a large dependence we found on measurement error due to the varying angle of incidence between antennas (Fig. 6 of [14]) and antenna radiation patterns (Fig. 7 of [17]). By restricting to two dimensions, we can orient the radios such that the antenna radiation patterns have equal energy over angle of incidence.

A. Problem statement

We start with a set of UWB nodes that are scattered through an environment. The size of this environment and what kinds of LOS or NLOS conditions to expect are unknown. We can attempt to take as many range measurements between any two nodes as we like, but depending on the distance between nodes and environmental conditions, it might not be possible to take measurements between every pair. Our goal is to find the coordinates and antenna delays of all the nodes given the set of measurements that we are able to make.

We pose this as a maximum a posteriori problem, expanding the prior on the parameters to include the probability of being able to measure a certain distance as well as that of encountering a certain antenna delay time:

$$\theta_{\text{opt}} = \arg \max_{\theta} p(\theta|M) \quad (1)$$

$$= \arg \max_{\theta} p(M|\theta)p(\theta) \quad (2)$$

$$= \arg \max_{\theta} \prod_{ij \in E} \left(p(\|X_i - X_j\|) p(M_{ij}|\theta) \right) \times \prod_i p(\text{delay}_i) \quad (3)$$

Optimization parameter θ includes both coordinates X and antenna delays; M includes all measurements taken on each edge in E (see symbol key in Table I).

We face a variety of problems making this formulation tractable. First, the individual measurements of any edge are not independent because they are conditioned on an unknown environment. For example, while LOS conditions generally result in a unimodal Gaussian distribution, NLOS conditions may lead to an arbitrary number of arbitrarily

TABLE I
SYMBOL KEY

c	speed of light, in m/ μ s
M	set of all measurements t for all edges ij
E	set of all measured edges ij in the network
X	set of all estimated coordinates (2 or 3-D) X_i
f	Gaussian probability density function
delay	set of all antenna delays $delay_i$
μ_{delay}	nominal antenna delay value
σ_{delay}	nominal antenna delay standard deviation
θ	optimization parameters $\{X, \text{delay}\}$
tof_{ij}	time of flight for edge ij
prop_{ij}	propagation time estimate for edge ij
d_{ij}	distance estimate for edge ij
$G(M_{ij})$	convolution-estimated PDF for edge ij
pdf_{ij}	table of estimated PDF for edge ij

spaced Gaussians that may not even include one centered on the true distance, depending on the environment.

Second, the need to perform computations on individual measurements may pose computational problems when there are many measurements. Third, finding the correct prior probabilities for distances that can be measured (since the hardware devices have limits to their range) and for the antenna delays also increases the complexity of the problem.

Without further constraints on the coordinates and network, the maximal solution will not be unique. At the very least, the coordinate system requires a fixed origin and coordinate axes. This leaves a solution that should differ from the ground truth by at most one translation, rotation, and mirroring operation. However, especially in NLOS conditions, node connectivity might be low enough and uncertainty high enough that it is possible for a network topology to emerge with flip ambiguities that result in multiple optimal solutions [4] [3]. A *flip ambiguity* is when a network has multiple unique exact solutions. The simplest example occurs when a node has connectivity of only two and its position could be on either side of the two connected nodes. Even if the solution is unique, conditions close to a flip ambiguity may contribute to various local minima, so a good initial estimate will be necessary to find the best solution.

One of the key ideas in this paper is to build a probabilistic model of range measurements in terms of a density estimation problem, viewing each of these individual measurements as data that informs the model. Our approach is to essentially convolve each measurement with a Gaussian corresponding to the intra-modal error (which we characterize under controlled testing conditions before the fact) and then compute the sum over the Gaussians. For simplicity, we also disregard the prior probability for distances on measured edges, $p(\|X_i - X_j\|)$, as it is likely uniform in the basin of convergence of the optimal solution.

We take $p(M_{ij}|\theta)$ to be the sum of the measurement probabilities, with f the Gaussian PDF, M_{ij}^t the measurement t of edge ij , $\text{prop}_{ij}(\theta)$ the propagation time estimated from parameters θ including coordinates X and antenna delays, and σ^2 the nominal intra-modal measurement variance:

$$p(M_{ij}|\theta) \approx \sum_t f(M_{ij}^t; \text{prop}_{ij}(\theta), \sigma^2) \\ := G(M_{ij}, \text{prop}_{ij}(\theta)) \quad (4)$$

We recognize that this is essentially a convolution between M_{ij} and a Gaussian with $\text{prop}_{ij}(\theta)$ specifying the location of the result in that convolution. This means we only need to calculate that convolution once in the method's initialization. We treat G as an estimate of the probability density function (PDF) of M_{ij} . This operation is effectively the same as calculating the kernel density estimate with a Gaussian kernel.

This leaves us with our final form of the optimization problem:

$$\theta_{\text{opt}} = \arg \max_{\theta} F(\theta) \quad (5)$$

$$F(\theta) = \prod_{ij \in E} G(M_{ij}, \text{prop}_{ij}(\theta)) \prod_i p(\text{delay}_i) \quad (6)$$

To solve the problem of the arbitrary origin, rotation, and mirroring, we make use of the fact that our intended application is robot localization. Our algorithm can first assign these arbitrarily during triangulation. Later, we can have the robot drive in a small loop, redefining the origin as the robot's initial location and the x-axis as the direction of its first motion. We can then resolve the mirroring ambiguity by coercing the winding order of the robot's estimated trajectory to match the known winding order of the path.

To solve both the problem of flip ambiguities and of a good initial estimate, we make use of trilateration with robust triangles [4]. This method will also identify whether the network is such that flip ambiguities cannot be distinguished. We run the robust trilateration with a series of different initial measurement values drawn from the distribution G of each edge. This gives us a variety of initial conditions across different possible modes of the data. We choose the trilateration with the lowest residual to further optimize.

We also repeat the entire process of choosing an initial estimate and further optimizing it several times to choose the best solution (see Algorithm 1).

B. Estimating probability density functions

We approximate the PDF of each link in the network by convolving the measurements with a Gaussian of nominal standard deviation $\sigma = 1.3 \times 10^{-4} \mu\text{s}$ (see Figure 2). The convolution is stored in a table over a range from $\min_j p_j - 6\sigma$ to $\max_j p_j + 6\sigma$ with a spacing of $2.5 \times 10^{-6} \mu\text{s}$, or approximately 0.075 cm.

We estimate this nominal standard deviation empirically as the standard deviation of measurements between a pair of our nodes under ideal unimodal LOS conditions.

We use quadratic interpolation to read out probabilities from each table. For queries beyond the 6σ range, we approximate the result with a Gaussian centered at the mean of the distribution with a standard deviation of $\sigma = 1.3 \times 10^{-4} \mu\text{s}$ again.

C. Robust quadrilateral trilateration

We find our initial localization estimate by assuming a fixed antenna delay value (μ_{delay}) for every node and fixed propagation times drawn from the measurement PDFs.

Algorithm 1:

Input: number of total attempts A
number of random initializations L
measurements M
edges E

Output: best solution θ_{opt}

```

for  $ij \in E$  do
  | pdf $_{ij} \leftarrow \text{EstimatePDF}(M_{ij})$ 
end
delay $^0 \leftarrow \{\mu_{\text{delay}}, \mu_{\text{delay}}, \dots\}$ 
for  $a \leftarrow 1$  to  $A$  do
  | for  $l \leftarrow 1$  to  $L$  do
    | prop  $\leftarrow \text{DrawFromPDFs}(\text{pdf})$ 
    |  $d \leftarrow \text{EstimateDistances}(\text{prop}, \text{delay}^0)$ 
    |  $X^{0l} \leftarrow \text{RobustTrilateration}(d)$ 
  | end
  |  $X^0 \leftarrow \arg \min_l \text{Residual}(X^{0l}, \text{delay}^0)$ 
  |  $\theta_a \leftarrow \arg \min_{\theta} F(\theta)$ 
  |  $\theta^0 = \{X^0, \text{delay}^0\}$ 
end
 $\theta_{\text{opt}} = \arg \min_a \text{Residual}(\theta_a)$ 

```

With these assumptions, we have a single distance measurement for each link to use in the trilateration:

$$\text{tof}_{ij} = \text{prop}_{ij} - 0.5(\text{delay}_i + \text{delay}_j) \quad (7)$$

$$= \text{prop}_{ij} - \mu_{\text{delay}} \quad (8)$$

$$d_{ij} = c \times \text{tof}_{ij} - \text{bias}(c \times \text{tof}_{ij}) \quad (9)$$

The bias function is a range bias characterized by the manufacturer that we have smoothed with quadratic interpolation [18]. From this point, the method is performed as described by Moore et. al. [4].

D. Minimization of the non-parametric objective function

In order to capture the full complexity of our measurements, we are building non-parametric error models of each link at run time. As this precludes use of analytic gradients, we optimize $\log F(\theta)$ with the generalized Gauss-Newton method and calculate both the gradient and Hessian numerically.

We use optimization parameter $\theta = \{X, \text{delay}\}$, where X includes both x and y coordinates. Extension to higher dimensions is trivial at this point.

$$\log F(\theta) = \sum_{ij \in E} -\log G(M_{ij}, \text{prop}_{ij}(\theta)) + \sum_i -\log p(\text{delay}_i) \quad (10)$$

$$\text{prop}_{ij}(\theta) = \text{tof}_{ij}(\theta) + 0.5(\text{delay}_i + \text{delay}_j) \quad (11)$$

$$\text{tof}_{ij}(\theta) = (d_{ij} + \text{bias}(d_{ij})) / c \quad (12)$$

$$d_{ij} = \|X_i - X_j\| \quad (13)$$

$$p(\text{delay}_i) = f(\text{delay}_i; \mu_{\text{delay}}, \sigma_{\text{delay}}^2) \quad (14)$$

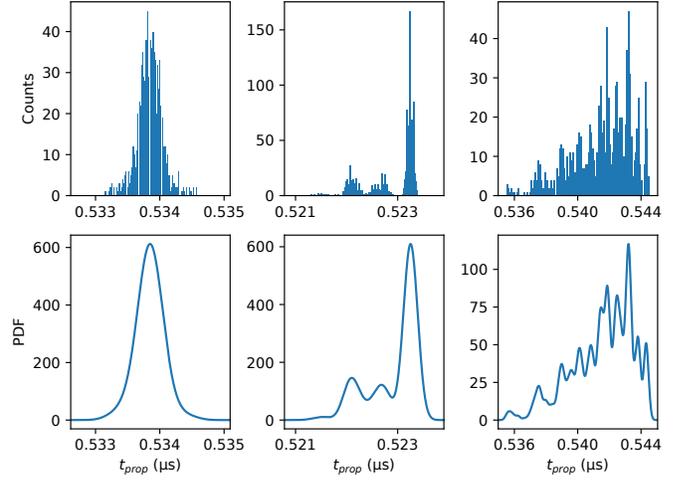


Fig. 2. Estimated probability distributions (*Bottom*) from 1024 UWB node-to-node measurements (*Top*) with varying line-of-sight conditions. A unimodal line-of-sight link (*Left*) could be easily approximated by a single Gaussian and the multimodal measurements (*Center*) could be fairly well represented by 3 or 4 Gaussians, but the heavily non-line-of-sight measurements (*Right*) are not Gaussian. Our non-parametric error model preserves the distribution in all three cases.

We use $\sigma_{\text{delay}} = 3.3 \times 10^{-4} \mu\text{s}$, which is derived from the manufacturer's antenna delay calibration note citing a 3-sigma variation of 30 cm[1]. Unfortunately, the manufacturer did not include a mean antenna delay, so we performed the manual calibration described in that document for 8 of our UWB nodes, and took $\mu_{\text{delay}} = 0.516 \mu\text{s}$ as the mean.

IV. EVALUATION

In all of our experiments, we configured our DWM1000 module to transmit on channel 1 at 850kbps with a 64MHz pulse repetition frequency (PRF), 512 symbol preamble, the DecaWave-optimized start-of-frame delimiter (SFD), and with smart transmit power control, according to the DW1000 user manual [19]. We used NLOS-optimized threshold settings, LDE_CFG1: 0x67, LDE_CFG2: 0x0004 [2].

For each dataset, we placed eight UWB nodes in our lab space as in Figure 3, with the antennas oriented vertically. As we intended to later localize the robot by use of the seven fixed nodes, we did not collect a ground truth location for the robot. However, measurements between the robot and the other seven nodes were still included in the graph realization. Dataset 1 was taken as shown in Figure 3 (*Top*), so that LOS would exist between all nodes. Dataset 3 was taken as shown in 3 (*Bottom*). Datasets 2 and 4 had fewer and greater obstructions than dataset 3, respectively.

While we perform these experiments with known Z coordinates (mostly zero), the method should in principle extend to higher dimensions. However, a full 3D solution is inherently more sensitive to the geometry of the network and the number of measurements.

To take data for evaluation of our method, we attempted to collect 1024 measurements for each of the 28 possible links. As some NLOS links might be completely blocked,

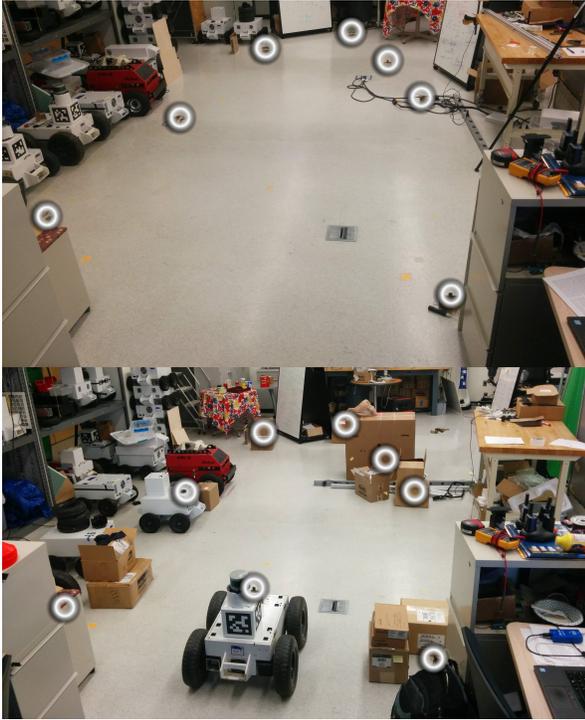


Fig. 3. Layout of our 7 fixed-location UWB radios, in line-of-sight conditions dataset 1 (Top), and non-line-of-sight dataset 3 (Bottom). Nodes are circled, but are only approximate in the bottom image. We made our datasets progressively non-line-of-sight by surrounding the nodes with boxes and other clutter.

we would give up if we could not collect that many measurements in under 10 seconds. Links with LOS conditions typically collected 1024 measurements in about 7 seconds. For all of our evaluations, we first process the data slightly to remove any measurements of less than $0.5 \mu\text{s}$ as obviously representing invalid negative ranges, and any measurements more than 4 standard deviations from the mean. This prevents the PDF tables from becoming needlessly large in an attempt to include outlier noise. In most of our evaluations, we do not use the full number of measurements taken, and so we take a random bootstrap sample of the desired number of measurements. We evaluate each condition with the same 50 distinct random seeds, so that each configuration is compared against the same 50 random subsets of measurements. We report the root mean square error between solved positions of the 7 fixed non-robot nodes and ground truth after finding the translation, rotation, and mirroring that minimizes this error. Ground truth was determined by placing nodes only at the corners of 1ft-by-1ft floor tiles.

Unless otherwise specified, we performed all of our evaluations with total method attempts $A = 20$, random trilateration initializations $L = 250$, $\mu_{\text{delay}} = 0.516 \mu\text{s}$, $\sigma_{\text{delay}} = 3.3 \times 10^{-4} \mu\text{s}$, and a maximum of 256 measurements per link.

We did not further investigate A or L because we found those values to be relatively insensitive, and performance did not seem to increase for larger values. We performed parameter sweeps over the variables describing the prior on antenna delay due to their significance to the method, and over the

number of measurements because taking measurements is the dominant factor in the total time taken for graph realization. The results are in Figures 4, 6, and 7.

Finally, we evaluate the performance of our method by ablation analysis. In the *first peak* version of the algorithm, we construct the PDF with a single Gaussian of mean μ centered on the lowest peak (local maxima) of the full PDF G , with standard deviation fit to minimize the least-squares error. This method follows the intuition that because incorrect multipath measurements are too long by definition, the first peak most likely represents the direct path. The *max peak* version is similar, but uses the highest amplitude peak of the full PDF, following the intuition that multipath or noise-based peaks will be lower in strength. The rest of the algorithm in both methods is unchanged. In the *triangulation* version, the algorithm skips the non-linear minimization of F (antenna delays remain at nominal values). In the *fixed delays* version, the non-linear minimization is only over the coordinates, leaving the antenna delay times constant at the prior mean value. The *manual delays* version is similar, but we manually calibrate the values in LOS conditions with LOS optimized thresholds [2]. We place the nodes three at a time in an equilateral triangle with 10-foot long sides and take 1000 measurements between each pair. We use the mean values and compute the least squares solution for the antenna delays [1]. The results of the total ablation analysis are in Figure 1.

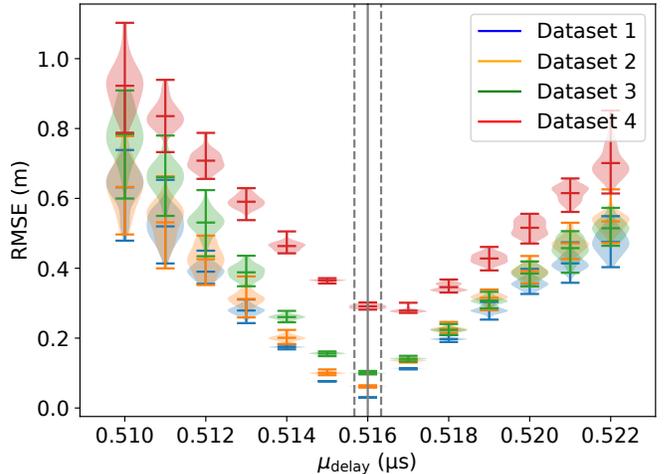


Fig. 4. Graceful degradation of our method’s accuracy when supplied with an incorrect antenna delay prior mean μ_{delay} . Each additional nanosecond difference from the empirical mean of $0.516 \mu\text{s}$ (vertical line) is a 30 cm biased error on each measurement. The DWM1000 manufacturer specifies an antenna delay standard deviation of $3.3 \times 10^{-4} \mu\text{s}$ (dashed lines) [1].

V. DISCUSSION

We use the ablation analysis in Figure 1 to examine the performance of our method under the increasing NLOS conditions from datasets 1 through 4. We find that under the most LOS conditions, the more naive first/max peaks methods have equivalent performance, but the disparity grows rapidly as the conditions become increasingly NLOS. We also find that the benefit of modeling antenna delays seems to fall

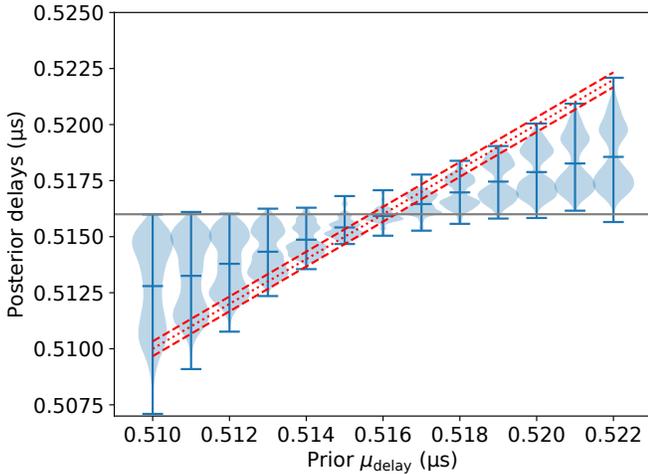


Fig. 5. Robustness of our method’s ability to solve for antenna delays with an incorrect prior mean (aggregated over all UWB nodes and datasets). If our method did nothing to improve on the given prior, each violin plot would have its posterior mean centered on the given prior (*dotted red line*). If our method were perfect, each plot would have the posterior centered on the known true mean (*solid horizontal line*). Because the prior standard deviation is small (*dashed red lines*), the optimization faces a high cost to move the posteriors away from the given prior. We see that under all conditions, our method produces posterior antenna delays significantly better than the supplied prior.

off under NLOS conditions. We recognize that with only 8 nodes, our complete problem is only barely fully constrained, having $8 \times 7/2 = 28$ equations and $8 \times 3 = 24$ unknowns. As the uncertainty in the NLOS links increases, the overall system would need to be more over-determined to accurately solve for all the unknowns. As a result, in high NLOS cases we cannot solve for the antenna delays to a higher accuracy than the $0.516 \mu\text{s}$ prior. However, on the most well constrained LOS case, modeling of the antenna delays cuts RMSE from 6.2 cm to 3.0 cm.

Examining the sensitivity of our method to the prior on the antenna delays, we find from Figure 4 that the prior mean μ_{delay} is significant to the performance of our method, but looking at Figure 5 we see that our method is able to return antenna delay values that are many standard deviations more accurate than the values provided by the incorrect priors. Looking at the prior on standard deviation σ_{delay} , we find the method to be far less sensitive. In the most NLOS case where the antenna delays are least constrained, we find that the too-large standard deviations are particularly harmful. Values in the range of about $1 \times 10^{-4} \mu\text{s}$ to $5 \times 10^{-4} \mu\text{s}$ seem to work best. Smaller values effectively prevent the antenna delays from moving away from the mean, reducing the algorithm to the fixed delay simplification, while larger values leave the choice of antenna delay too unconstrained.

The total time necessary for the graph realization and calibration in a new environment is dominated by the time needed to take measurements across all possible edges. With our evaluation setup, taking 128 measurements per link would take a maximum of 35 seconds for the 8 nodes. Execution of the complete algorithm took half a second on the author’s laptop in the worst case, leading to less than one minute for the complete localization process.

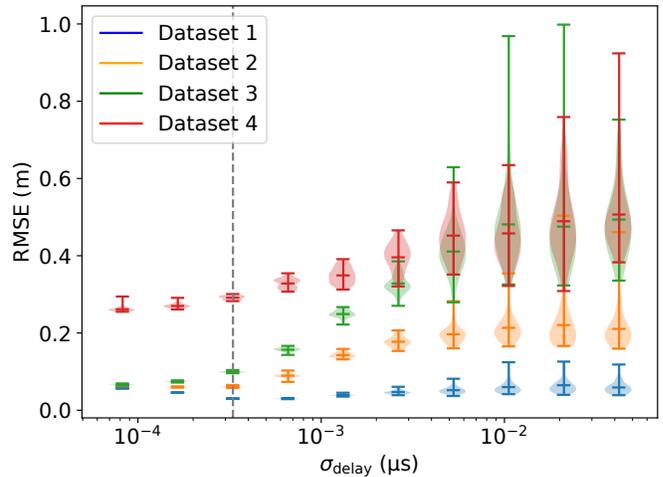


Fig. 6. Our method has low sensitivity to antenna delay standard deviation prior σ_{delay} . Dataset 1 has low uncertainty and so is well constrained, making it insensitive to choice of σ_{delay} . Datasets 2, 3, and 4 have increasing uncertainty from non-line-of-sight conditions, making them more sensitive. Values below the nominal of $3.3 \times 10^{-4} \mu\text{s}$ (*dashed vertical line*) essentially make the method perform as if the antenna delays are fixed. Datasets 3 and 4 are not quite constrained enough to solve for antenna delays well, so their RMSE grows the most for larger σ_{delay} (compare the full method with the fixed antenna delay version in Figure 1).

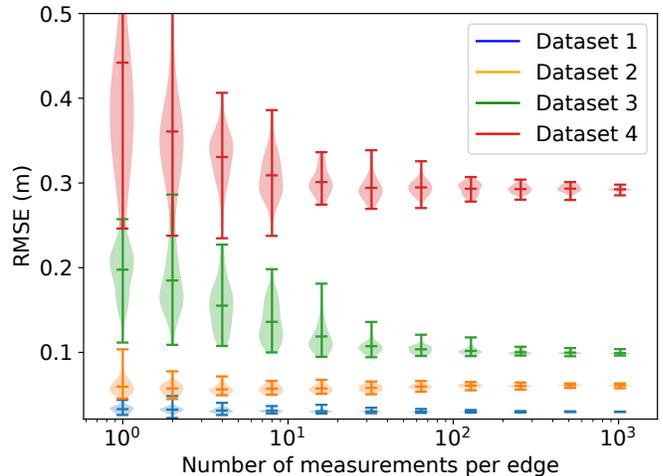


Fig. 7. Convergence rate of our method on varying line-of-sight conditions. The fewer the measurements needed, the less time required to localize and characterize the network. In datasets 3 and 4, accuracy continues to improve until about 128 measurements per edge. In datasets 1 and 2, however, low measurement uncertainty allows even just one measurement per link to suffice.

VI. CONCLUSION

In summary, we have found that our non-parametric error modeling makes our method robust to NLOS conditions, and our modeling and simultaneous calibration of antenna delays leads to very high accuracy in LOS conditions. We found that only 16 to 128 measurements are necessary per link and that a network of 8 UWB nodes requires less than a minute for both the collection of measurements and execution of the algorithm. On this network we achieved accuracies of 3 cm RMSE in the LOS dataset 1, and 30 cm in the NLOS dataset 4. Our datasets and source code files can be found at <https://osf.io/pkbq4/>.

REFERENCES

- [1] *APS014: DW1000 Antenna Delay Calibration*, decaWave, 2014, version 1.01. [Online]. Available: <https://decawave.com/sites/default/files/aps014-antennadelayscalecalibrationofdw1000-basedproductsandsystems.v1.01.pdf>
- [2] *APS006 Part 2 Application Node: Non Line of Sight Operation and Optimizations to Improve Performance in DWM1000 Based Systems*, decaWave, 2014, version 1.4. [Online]. Available: https://decawave.com/sites/default/files/aps006_part2_nlos_operation_and_optimizations.pdf
- [3] A. A. Kannan, B. Fidan, G. Mao, and B. D. Anderson, "Analysis of flip ambiguities in distributed network localization," in *Information, Decision and Control, 2007. IDC'07.* IEEE, 2007, pp. 193–198.
- [4] D. Moore, J. Leonard, D. Rus, and S. Teller, "Robust distributed network localization with noisy range measurements," in *Proceedings of the 2nd international conference on Embedded networked sensor systems.* ACM, 2004, pp. 50–61.
- [5] K. Chintalapudi, A. Padmanabha Iyer, and V. N. Padmanabhan, "Indoor localization without the pain," in *Proceedings of the sixteenth annual international conference on Mobile computing and networking.* ACM, 2010, pp. 173–184.
- [6] R. Faragher and R. Harle, "Location fingerprinting with bluetooth low energy beacons," *IEEE journal on Selected Areas in Communications*, vol. 33, no. 11, pp. 2418–2428, 2015.
- [7] J. Wang, Q. Gao, Y. Yu, H. Wang, and M. Jin, "Toward robust indoor localization based on bayesian filter using chirp-spread-spectrum ranging," *IEEE Transactions on Industrial Electronics*, vol. 59, no. 3, pp. 1622–1629, 2012.
- [8] C. Zhang, M. Kuhn, B. Merkl, M. Mahfouz, and A. E. Fathy, "Development of an uwb indoor 3d positioning radar with millimeter accuracy," in *Microwave Symposium Digest, 2006. IEEE MTT-S International.* IEEE, 2006, pp. 106–109.
- [9] A. A. Kannan, G. Mao, and B. Vucetic, "Simulated annealing based localization in wireless sensor network," in *Local Computer Networks, 2005. 30th Anniversary. The IEEE Conference on.* IEEE, 2005, pp. 2–pp.
- [10] T.-C. Liang, T.-C. Wang, and Y. Ye, "A gradient search method to round the semidefinite programming relaxation solution for ad hoc wireless sensor network localization," *Sanford University, formal report*, vol. 5, 2004.
- [11] X. Ji and H. Zha, "Sensor positioning in wireless ad-hoc sensor networks using multidimensional scaling," in *INFOCOM 2004. Twenty-third Annual Joint Conference of the IEEE Computer and Communications Societies*, vol. 4. IEEE, 2004, pp. 2652–2661.
- [12] C. Di Franco, A. Prorok, N. Atanasov, B. Kempke, P. Dutta, V. Kumar, and G. J. Pappas, "Calibration-free network localization using non-line-of-sight ultra-wideband measurements," in *Proceedings of the 16th ACM/IEEE International Conference on Information Processing in Sensor Networks*, ser. IPSN '17. New York, NY, USA: ACM, 2017, pp. 235–246. [Online]. Available: <http://doi.acm.org/10.1145/3055031.3055091>
- [13] B. Kempke, P. Pannuto, B. Campbell, and P. Dutta, "Surepoint: Exploiting ultra wideband flooding and diversity to provide robust, scalable, high-fidelity indoor localization," in *Proceedings of the 14th ACM Conference on Embedded Network Sensor Systems CD-ROM.* ACM, 2016, pp. 137–149.
- [14] F. Despoux, K. Jaffres-Runser, A. Van den Bossche, and T. Val, "Accurate and platform-agnostic time-of-flight estimation in ultra-wide band," in *Personal, Indoor, and Mobile Radio Communications (PIMRC), 2016 IEEE 27th Annual International Symposium on.* IEEE, 2016, pp. 1–7.
- [15] J. D. Hol, F. Dijkstra, H. Luinge, and T. B. Schon, "Tightly coupled uwb/imu pose estimation," in *2009 IEEE International Conference on Ultra-Wideband.* IEEE, 2009, pp. 688–692.
- [16] L. Zwiello, X. Li, T. Zwick, C. Ascher, S. Werling, and G. F. Trommer, "Sensor data fusion in uwb-supported inertial navigation systems for indoor navigation," in *2013 IEEE International Conference on Robotics and Automation.* IEEE, 2013, pp. 3154–3159.
- [17] *DWM1000 Datasheet*, decaWave, 2016, version 1.6. [Online]. Available: <https://www.decawave.com/sites/default/files/resources/DWM1000-Datasheet-V1.6.pdf>
- [18] *APS014 Application Node: Sources of Error in DW1000 Based Two-Way Ranging (TWR) Schemes*, decaWave, 2014, version 1.0. [Online]. Available: https://decawave.com/sites/default/files/aps011_sources_of_error_in_twr.pdf
- [19] *DW1000 User Manual*, decaWave, 2017, version 2.15. [Online]. Available: https://www.decawave.com/sites/default/files/dw100020user20manual_0.pdf